

The Robot Baby and Massive Metacognition: Early Steps via Growing Neural Gas

Jared Shamwell¹, Tim Oates², Preeti Bhargava¹, Michael T. Cox¹, Uran Oh¹, Matthew Paisner¹, Don Perlis¹
¹University of Maryland, College Park ²University of Maryland, Baltimore County
Email: {shamwell,prbharga,mcox,uranoh,mpaisner,perlis}@cs.umd.edu, {oates@cs.umbc.edu}

Abstract— We have initiated a long-term robotics project based on our previous work on metacognition as a powerful tool that can synergistically play machine learning and commonsense reasoning off one another. The new project involves a mobile robot that lives in a room and learns about the room and about itself. The robot is initially set up to have a standard set of facilities (vision, IR, limb, wheels, planners, learning modules, some modest NLP, a reasoner, etc.) but it does not know much about its capabilities or how to properly use them. It has a prime directive: to learn. This paper will focus on one of the first major questions of this project: can we use Growing Neural Gas (GNG) to discover the physical structure of an environment and, if so, what are the limits of its use? To answer this question, we have devised an experiment to test whether our robot can distinguish between two identical objects using only GNG. Preliminary results suggest that passing image data along with robotic control signal data is sufficient to autonomously detect the basic physical structure of a room.

Index Terms—Growing neural gas, learning, robotics, sensorimotor, metacognition.

I. INTRODUCTION

Human infants learn about their environment, their bodies, and the relationship between the two from sensorimotor exploration over long periods of time. Robots, on the other hand, tend to have much of this knowledge provided a priori, and knowledge acquired from learning either does not impact future learning or is heavily curated. In this paper we explore ongoing unsupervised learning about the structure of the environment and the effects of a robot’s action, more akin to a human infant

This work is part of a long-term robotics project based on our previous work on metacognition as a powerful tool that can synergistically play machine learning and commonsense reasoning off one another. For a more complete overview of our project, see our companion paper [1] and for our previous work on metacognition, see [2, 3, 4, 5].

We have chosen to focus on one of the first major questions of this project: can we use Growing Neural Gas (GNG) to discover the physical structure of an environment and, if so, what are the limits of its use? In answering this question, we will begin to craft a vision of the robot’s architecture and which aspects can be built from a bottom-up, data-driven framework and those needing a top-down framework.

II. GROWING NEURAL GAS

Growing neural gas (GNG) (see Fritzsche [6]) is an unsupervised learning algorithm that is intended to give an agent the ability to autonomously create a model of states in its environment without having any prior knowledge of that environment. This formulation is extremely useful in the context of developmental robotics, because it allows our robot which has little to no a priori knowledge of its environment to access powerful tools such as planning or logical reasoning that require a description of discrete states and actions to perform optimally.

Wichert [7] proposed that it is also possible for a mobile robot to model the topology of the environment with visual perception with growing neural gas [6]. Baldassarri and colleagues [8] compared the two most well-known self-organizing networks, Self Organizing Maps (SOM) and GNG, on corridor learning tasks for a mobile robot with a visual sensor. They have shown that GNG performs better in preserving topology thanks to its dynamic adaptation.

III. LEARNING THE ENVIRONMENT

A. The Problem

We are interested in GNG’s ability to enable autonomous robotic learning using an array of raw sensory data (e.g., visual, IR, velocity/odometric, depth, etc.). At this juncture, we have chosen not to focus on the vector quantization of the image data as in Fritzsche [6], but have instead directed ourselves toward the effects of including multiple classes of disparate data types in the input vectors passed to GNG.

As a beginning, we are interested in whether we could use GNG to extract information about the physical structure of an environment from raw image data. Wichert has already shown that the physical location at which a robot captures an image can be reflected in the GNG cluster to which it belongs [7]. We wish to expand upon this work by extracting information about how areas of the environment are connected. Our ultimate aim is to see whether GNG and other techniques can be used effectively to sort out useful patterns in a wide range of data types including visual, tactile and even internal state data.

In learning the structure of a room, or several rooms, the robot will need to be able to reliably figure out that one object is different from another identical object. We endeavored to




Fig. 1. GNG with time data and multiplier of 100. Circled nodes indicate the two clusters each corresponding to one of the two different red objects.

determine whether this could be accomplished directly with GNG or if it had to be accomplished through a higher-level mechanism (e.g., our reasoner).

B. Experiment

For our investigation, we collected sets of images from a real robot moving around in a single room. The robot saved images at a rate of approximately 15 frames per second. We used our own java implementation of [6]’s GNG algorithm. For all of our experiments, GNG was run for 100,000 iterations on sequential 320x240 pixel images down sampled to 40x30 pixels using a simple pixel area averaging algorithm.

We had the robot spin in circles and captured a set of images when spinning clockwise at 0.2 r/s. Each image was also marked with the time at which it was taken. In each image set, images were saved while the robot spun in a circle for a single complete rotation. The first image in each set was marked with a zero and each subsequent image was incremented by one. This counter was used later as a time stamp. The room was set up with unicolor sheets on the walls and two identical red trash cans in its search path.

We began by determining the optimal network sizes for our qualitative analysis of the GNG networks generated from our image data. With our sets of images, a maximum network size of 20 results in graphs that exhibit levels of detail sufficient to visually/qualitatively extract major patterns without excess confusion. While the algorithm enables GNG to decide its own maximum node count based on a given metric (e.g., average node error), we wished to use a standard static node count for all of our experiments as the use of any such metric would result in different node counts depending on the experiment.

To investigate the problem of differentiating between identical objects with GNG, we extended the input vector to contain the time tick on which each image was collected. We discovered that passing the temporal data to GNG as a single non-normalized value in the input vector alongside the image data does not result in a graph that is significantly different from the graph generated from the same images without temporal information.

This is likely due to (1) the number of elements in the input vector corresponding to image data dwarfing the single value

corresponding to time; and (2) the lower variance in the time values added to the input vector compared to the variance in the image data places a greater influence on image data.

We next used a static multiplier on the time values passed added to the input vectors to overcome these difficulties and found that using a multiplier of 100 best captures the structure of the room for this particular dataset and, as can be seen in Fig 1., results in two clusters of red objects corresponding to our two red objects.

IV. CONCLUSION AND FUTURE WORK

Preliminary results suggest that directly passing raw image data accompanied with robotic control signal data to GNG is sufficient for detecting the basic physical structure of a room. In these experiments, our approach to detecting multiple identical objects relies on images being marked according to a closed continuous timespan. As we continue to use the robot in other applications we will not have this luxury, which leaves this approach non-scalable. In response, we are planning to include localization/odometric data in place of the temporal data when passing information to GNG. We believe that this is an important next step in leveraging GNG to detect the structure of an environment and the spatial relations of objects therein. Furthermore, our experiments have demonstrated GNG’s potential for use in more complex robotic reasoning and learning tasks involving disparate data streams. GNG provides the means for a data-driven approach to learning with a top-down mediator in MCL.

REFERENCES

- [1] Bhargava, P., Cox, M., Oates, T., Oh, U., Paisner, M., Perlis, D., and Shammwell, J. (in press). “The Robot Baby and Massive Metacognition: Future Vision.” 2012 IEEE International Conference on Development and Learning.
- [2] Schmill, M., Anderson, M., Fults, S., Josyula, D., Oates, T., Perlis, D., Haidarian, H., Wilson, S., and Wright, D. (2011). The Metacognitive Loop and reasoning about anomalies. In Cox, M., Raja, A. (Ed.), *Metareasoning: Thinking about thinking* (pp. 183-198). Cambridge, MA: MIT Press.
- [3] Josyula, D.P., Vadali, H., Donahue, B.J. and Hughes, F.C. (2009). “Modeling Metacognition for learning in Artificial Systems.”, In Proceedings of the International Symposium on Innovations in Computing (INC’09).
- [4] Josyula, D.P., Fults, S., Anderson, M.L, Wilson, S., Perlis, D. (2007). “Application of MCL in a dialog agent”, In Papers from the Third Language and Technology Conference.
- [5] Josyula, D.P., Donahue B., McCaslin M., Snowden M., Anderson M., Schmill, M., Oates, T., and Perlis, D. (2010). “Metacognition for Detecting and Resolving Conflicts in Operational Policies”, Proceedings of the Workshop on Metacognition for Robust Social Systems at the 24th AAAI Conference on Artificial Intelligence (AAAI’10).
- [6] Fritzke, B. (1995). “A Growing Neural Gas Networks Learns Topologies”, *Advances in Neural Information Processing Systems*, 7, Brussels 625-632.
- [7] von Wichert, G. (1996). “Selforganizing Visual Perception for Mobile Robot Navigation”, In Proceedings of EUROBOT.
- [8] Baldassarri, P., Puliti, P., Montesanto, A., and Tascini, G. (2003). “Self-Organizing Maps versus Growing Neural Gas in a Robotic Application”. In *Artificial Neural Nets Problem Solving Methods: 7th International Work-Conference on Artificial and Natural Neural Networks, IWANN2003, Part II, Lecture Notes in Computer Science*.